Heating/Cooling Notes

Leon Mualem

University of Minnesota

This note documents some data and calculations that I made with regard to the prototype APD readout board and cooling system. It also has some application to the final APD readout box which will interface to modules in the NOvA detector.

Heat sources

In the current design there are a few heat sources whose power the TE cooler must remove from the box. Heat can flow into the APD via conduction along the fiber (or any connector we choose to use), it can be conducted from the PCB through the pins (or flip chip bonds), and it can be delivered via the air around the APD.

Thermal power is determined by the thermal conductivity of the material, the area of the material, the length of the material, and the temperature difference.

$$\dot{Q} = \frac{kA}{L}\left(T_1 - T_2\right)$$

Fiber

Polystyrene has thermal conductivity, k = 0.125W/m/K. For the baseline detector the fibers are 0.8mm, and there are approximately 128 fiber ends per box. The total area is:

$$A = 128\pi\left(\frac{0.8mm}{2}\right)^2 = 64\mu m^2$$

If you assume that the total fiber length from the outside (assumed to be ambient temp, 25C), and the APD surface (assumed to be at baseline temp of –15C) is only 1 *cm* then the total power delivered via the fiber is thus:

$$\dot{Q} = \frac{0.125W/m/K * 64\mu m^2}{1cm}\left(40^\circ C\right) = 32mW$$

In fact, this is probably an upper limit on the power delivered by the fiber, as any reasonable fiber connection will probably be longer than this, thus reducing the power from the fiber itself.

If a connector interfaces directly to the APD, it will have area of about 5mmX15mm, or for an entire box 4 times this area or 0.3E-3m^2, or about 5 times the area of the fiber. Even so, the heat load would be 160mW, which is still quite small. So assuming the connector is made of a plastic similar to polystyrene, the heat load should not be significant from this source.

APD Heating

The power of the APD itself is very small. The typical power dissipated by the device itself is due to the dark current and the operating voltage. The typical dark current in one of the diode arrays is at most 4nA/channel, totaling 64nA. The typical operating voltage is 400V, so the power is 64nA*400V=25μW. Thus the total power in a box due to the APD power is only 0.1mW.

PCB

Generally the heat flow to the APD will be limited by the PCB. The pins, or contacts will not represent a significant thermal resistance since they are short and highly conductive in any case. Therefore all the thermal power will come from devices on the board, or conduction within the board from the ambient temperature.

As an example, consider the heat load from the MASDA chip, the dominant heat source on the APD prototype board, and the one closest to the APD itself. If we assume that the full cross-sectional area of the board is the conductive cross section, then the area is 1.5mm x 75mm, the thickness and width of the board. The distance from the MASDA is about 20mm. If we assume the FR4, as a fiberglass/epoxy sandwich has a thermal conductivity of 0.2, typical of a plastic, and rather high for a glass, and that the temperature difference is about 50°C the thermal power delivered is:

$$\dot{Q} = \frac{0.2 \times 1.5mm \times 75mm}{20mm} \times 50°C = 56mW$$

In fact, the MASDA is dissipating nearly 0.5W of power, which mainly dissipates by heating the air. (Of course this power needs to be removed too, but it can be done with a small fan, or heat sink radiator in the “warm” side of the box.

Air

In fact, the dominant heat flow into the prototype system is via conduction through the metal box and the air inside it. It was measured during the prototype testing that the temperature inside the box at the APD operating temperature of –15°C and cooling water temperature of 4°C that the temperature inside the box was 10°C lower than outside,16 and 26°C respectively.

The conductivity of aluminum is quite high, 205W/m·K, and the surface area of the box high also, approximately 1320cm^2. With a thin skin of about 1mm the amount of heat that can flow for each °C difference is 24kW! Thus this is not a limitation in heat flow into the box. The air inside the box is not as conductive though. Air itself has a thermal conductivity of only 0.024W/m·K. However, since air is not opaque to infra-red radiation, and it is allowed to move its effective conductivity is much greater. To get some estimate this I use the R-value, yes the insulation one. Using the R-value/in it is possible to derive an effective conductivity of air, and thus estimate the heat load imposed by the conductivity of the air. I determined that the conductivity assuming an air R-value of 1/3.5” is 0.5W/m·K. Thus the heat load is

$\dot{Q} = \frac{0.5(W/m \cdot K) \times 1320(cm^2)}{2.5(cm)} = 2.6W/K$. For a 10°C difference this is 26Watts. This is probably an overestimate, as much of this power is probably removed by the chilled water heat sink, and can probably be eliminated by reducing the cold area, and insulating the cold section of the box. Typical foam insulations have 1/10 the thermal conductivity assumed in this section. By injecting some foam inside plastic it would be possible to insulate the box and retain serviceability.

Aside on R-Value

R-Value is simply the thickness of the material divided by the thermal conductivity:

$$R = \Delta x / k$$

Thus the R-value increases with thickness and decreases as thermal conductivity increases. Thus using the R-value of a material and its thickness, or R-value/inch or cm it is possible to determine the thermal conductivity. This can also take into account some of the other effects of conduction, convection and radiation which are all rolled up together, as the R-value of a material.

The typical R-value unit is derived from the temperature difference in degrees Fahrenheit, the area in square feet, the time in hours, and the heat loss measured in BTUs:

$$R = \frac{\Delta T(°F) \times A(ft^2) \times t(hr)}{Q(BTU)}$$

To convert these units into those useful for scientists, the most useful unit is the R-value/in. Given an R-value/in the thermal conductivity is only a simple conversion away.

The thermal conductivity $k = 0.144 * \frac{1}{R/in}(W/m \cdot K)$.

Cooling Power in Prototype

In actual experience the cooling power that has been required to keep our cookie/APD system cooled to –15C is about 3.5W. This is in good agreement with the temperature difference of 19°C and ~3W of power delivered to the TE cooler module. The total power in the cooling system is somewhat higher, but depends on what voltage the controller is using, but is approximately 6W. TE coolers are actually quite efficient if the temperature difference is small. For example, if you used an ice bath on the "hot" side of the TE cooler the temperature difference would be 15°C. For the cooler I am currently using this would correspond to an input electrical power of 2.8W and a cooling power of 3.3W, so you actually remove more heat than the electrical power you input. The relevant measure in heat pumping applications is the coefficient of performance (COP), which is the ratio of cooling power to input electrical power. In this case the coefficient of performance is about 1.2. If you have a temperature difference of 40°C you will have an input electrical power of about 7.5W for a cooling power of 3.25W, so you don't remove as much heat as electrical power, and your coefficient of performance is ~0.5. COP greater than 1 is not unusual for small temperature differences relative to the maximum temperature difference, which is typically about 70°C. For more information you might want to look at www.tetech.com, there is more than enough information there to fill up a afternoon.

The lowest temperature with the current setup was recorded with a total power to the TE module of 26W. The temperature was –27°C at the APD and 4°C at the water cooler.

The cooling power was only about 6.5W though since TE coolers get less efficient as you get closer to their limits. Here the COP is only 0.25.

Future Cooler

It is likely that we will need even less power to cool the APDs in the final board for at least 2 reasons. The first is that it will be possible to insulate the boxes and reduce the heat load. The second is that we will be cooling the APDs directly, instead of indirectly through the APD face as is required in the packaged version of the array. This will eliminate heat load derived from the cold cookie and cold-plate cooler itself. It is possible to have TE coolers designed to fit the particular application which will make the most efficient use of the coolers.

Some references

Insulation R-values
http://hyperphysics.phy-ast.gsu.edu/hbase/tables/rvalue.html
Thermal conductivity
http://hyperphysics.phy-astr.gsu.edu/hbase/tables/thrcn.html

Material	Conductivity (W/m·K)
Solid Polystyrene	0.125
Air	0.024
Aluminum	200
PVC	0.2
Water	0.6
Foam Polystyrene	0.025
Glass	0.8

Table 0.1 Some relevant thermal conductivities

Material	R-Value
Plywood	1.25/inch
Wood	1.25/inch
Foam Polystyrene	3.45/inch
Urethane foam	5.3/inch
Urea Formaldehyde foam	4.48/inch
Fiberglass board	4.5/inch
Glass	7.28/inch ? from 0.91 for 1/8" pane
Air gap (3.5")	1.01
Stagnant air layer	0.17

Table 0.2 Some R-values of potential interest

Other potentially useful numbers

Heat of fusion of water is 335kJ/kg. This means that one pound of ice will absorb 150kJ of heat as it melts, which can provide 42W·hr of cooling. So if you want to run one of the prototype systems with an ice bath as the heat sink you can do so for about 10 hours on one pound of ice. This assumes that all the heat going into the ice batch is from the cooler, so you would want an insulated ice bath, or you will need more ice.

Additional calculations for flip-chip bonding to PCB

Conductivity of solder balls for a bump bond

Each solder ball has a diameter of 135µm, and a height equal to 70% of the diameter. Due to the distortion of the ball as it wets to the surface, the area is more of a cylinder than a ball, so I will model the conductivity of the solder bump-bond as a cylinder with diameter of 135µm, and height of 100µm. The heat flow through this bond is then:

$\dot{Q} = \frac{67(W/m\cdot K)\times\pi\times(67.5\mu m)^2}{100(\mu m)} = 9.5mW/K$, and for an entire array there would be 50 of these capable of delivering 480mW/K. This conductivity is actually low enough, due to the small area of the bonds, that there can be a significant (~1degree) temperature difference between the APD array and the substrate.

Conductivity of an epoxy bond under the APD

A flip-chip bond is typically covered by an epoxy to relieve strain on the bond itself. The area of this epoxy in this design would be 4.4x15.1+10.6x1.2=79(mm)^2. This would have a thickness the same as the solder ball, or 100µm. The heat flow would then be:

$\dot{Q} = \frac{0.2(W/m\cdot K)\times 79(mm)^2}{100(\mu m)} = 0.158mW/K$, and therefore the heat would be almost entirely delivered by the solder balls, in spite of the much larger area of the epoxy, the much lower conductivity limits the heat that can be supplied via this path.

Heat flow along electrical connections (board traces)

For simplicity I will simply assume 2.54cm traces that are 17u high and 180u across. There will be at least 34 of these, maybe more for temp sensor, so I will assume 36.

$\dot{Q} = \frac{36 * 400.0(W/m\cdot K)\times 17\times 180(\mu m)^2}{2.54(cm)} = 1.7mW/K$, for a dT of 50C that is 85mW.

Heat load from fiber connector block (cookie)

The fiber connector should have an area of approximately 2cmX1cm, a length of at least 1cm, and be made of a plastic. I will assume a thermal conductivity of the plastic to be 0.33W/m/K. Therefore the heat flow through this cookie will be

$\dot{Q} = \frac{0.33(W/m\cdot K)\times 2.0(cm)^2}{1.0(cm)}(25-(-15)C) = 264mW$. It makes little difference in terms of heat flow if the fiber block is directly in contact with the APD or held off by 25-50µm. Even with the low conductivity of air, or dry nitrogen there is essentially no thermal resistance since the gap is so small. It will reduce the heat flow slightly, as the temperature will be reduced slightly. I suspect an even bigger effect will be what color the block is, as the IR radiation can be a significant part of the heat transfer.

Heat load from cooling the printed circuit board

There will be a heat load delivered by the PC board since the APD is mounted directly on the surface. The low thermal conductivity of FR4 (0.27W/m/K) helps to minimize this load. A significant reduction in conductivity from what one might expect comes from not having any ground or power plane in the area of the APD; the thermal conductivity of copper, even thin layers, would significantly alter the load.

In order to estimate what the load would be I did a brute force calculation of the steady state heat flow if the edges of a 3”x3” board were held at 25C, and a cold APD (-15C) was mounted in the center. This is not very unlike the current board, and any one might conceive. Using these boundary conditions I simply let heat flow from cell to cell in a grid and calculated the change in temperature. Once the changes had stopped (down to machine accuracy of ~1.5E-15C) the steady state solution was plotted, and the heat flow through the boundary of the cold cell was calculated. The total power delivered from the board is 67mW.

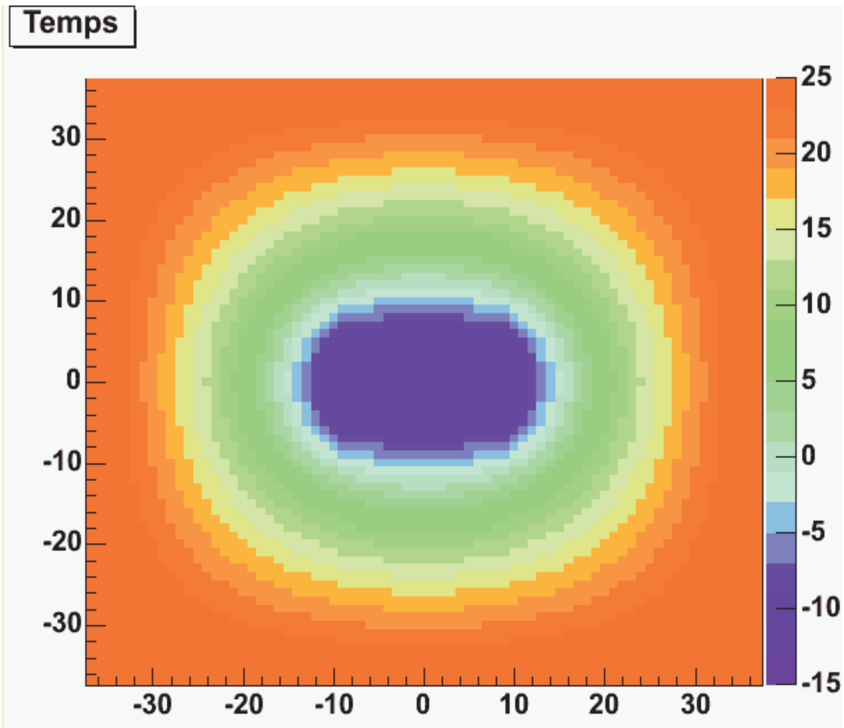

Figure 1 Steady state solution of board temperature for a 2cm by 1cm APD array held at –15C mounted at the center of a 3”x3” FR4 board. The edges of the board are held at 25C. The X and Y axes indicate position on the board (mm), and the color scale indicates temperature in Celsius at each position.

Changing some of the conditions, for example, adding a hot ASIC (70C) at the edge of the board does not significantly affect the heat load through the board (It does change the heat delivered by the board traces, the main heat load). This change added only 7mW to the heat delivered by the board. The steady state solution is shown in figure 2.

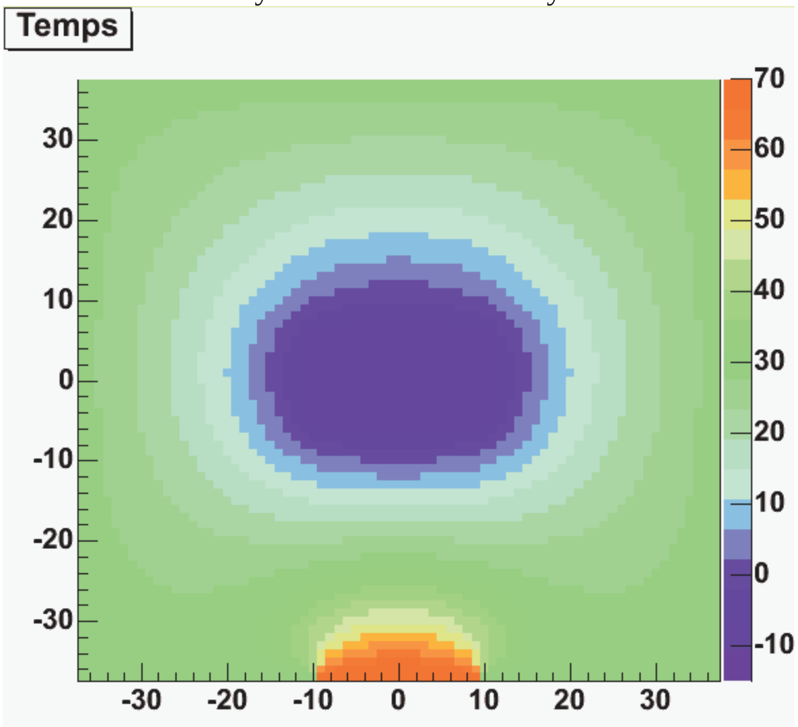

Figure 2 Steady state solution of board temperature for a 2cm by 1cm APD array mounted at the center of a 3"x3" FR4 board, the edges are held at 25C, except for the small red area, simulating an ASIC at 70C. The X and Y axes indicate position on the board (mm), and the color scale indicates temperature in Celsius at each position.

Moving the APD position close to the edge of the board does increase the heat load to 96mW, with a solution as shown in figure 3.

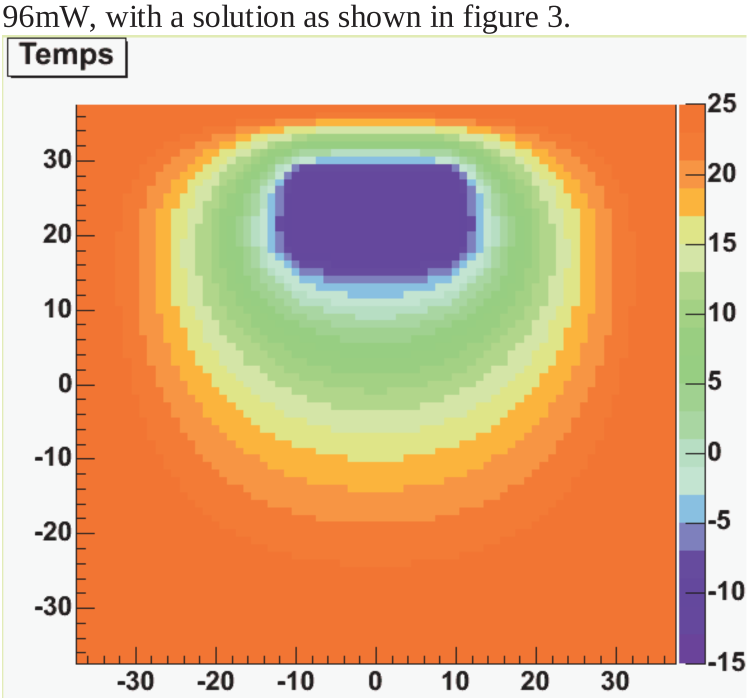

Figure 3 Steady state solution of board temperature for a 2cm by 1cm APD array held at –15C, and mounted 1cm from the edge of a 3”x3” FR4 board, the edges are held at 25C. The X and Y axes indicate position on the board (mm), and the color scale indicates temperature in Celsius at each position.

Total Heat Load (to TE cooler) Using Flip-Chip mounting

The total heat load that the TE cooler will have to deal with is the load from the fiber block, the traces on board, and the board itself.

$$\begin{aligned} \mathbf{Q(Total)} &= \mathbf{Q(Block)+Q(traces)+Q(PCB)} \\ &= \mathbf{264mW+85mW+67mW=416mW.} \end{aligned}$$

Reality Check

There is an additional dose of reality that should enter into this equation. (In fact, all of these equations.) Air is a lousy conductor. I assumed that the edge of the board is at 25C, and can supply as much heat as needed to keep it this way. In reality this is not how it works. The board cools down, and the heat load is LESS than calculated. Running the prototype board in the current system cools the air in the box by 10C. Essentially the air is an additional layer of insulation between the cold board and the infinite heat bath of the world.

There is also the fact that the heat (Q(PCB)+Q(traces)) will flow through the APD bump bonds. This will make the board temperature actually ½ degree higher than the APD temperature. This will only reduce the heat flow by about 1% though, so I neglected this, and a temperature gradient (from front to back) of the APD, ½, or 1/30th of a degree here or there is not going to make this design unworkable.

FINAL COMMENTS (in my dreams)

I like the idea of having the injection molded parts for the cookie, and the containment volume. I think we would have something along those lines whatever we do. It may be true that we need a containment volume on both the front (fibers coming in) side of the board, and the back, we don’t want cold parts exposed to “air” where condensation might take place.

I haven’t drawn pictures here, but in my mind they look a lot like those in Jon’s talk. The flip-chip just changes the order of things, moving the PC board to the top of the stack, as I indicated in my drawings that were included in the talk.